arXiv:1602.03600v1 [stat.ML] 11 Feb 2016

Data-Driven Online Decision Making with Costly Information Acquisition

Onur Atan OATAN@UCLA.EDU
University of California, Los Angeles

William Whoiles WHOILES@ECE.UBC.CA
University of California, Los Angeles

Mihaela van der Schaar MIHAELA@EE.UCLA.EDU
University of California, Los Angeles

Abstract

Existing work on online learning for decision making takes the information available as a given and focuses solely on choosing best actions given this information. Instead, in this paper, the decision maker needs to simultaneously learn both what decisions to make and what source(s) of contextual information to gather data from in order to inform its decisions such that its reward is maximized. We propose algorithms that obtain costly source(s) of contextual information over time, while simultaneously learning what actions to take based on the contextual information revealed by the selected source(s). We prove that our algorithms achieve regret that is logarithmic in time. We demonstrate the performance of our algorithms using a medical dataset. The proposed algorithm can be applied in many applications including clinical decision assist systems for medical diagnosis, recommender systems, actionable intelligence etc., where observing the complete information in every instance or consulting all the available sources to gather intelligence before making decisions is costly.

1. Introduction

In the classical contextual Multi-Armed Bandit (MAB) problem (Chu et al., 2011; Slivkins, 2011; Lu et al., 2010; Dudik et al., 2011; Langford & Zhang, 2007), a decision maker repeatedly observes some contextual information, selects an decision-action, and observes a random reward. The realized reward for each choice provides information about the distribution of rewards associated with each context and decision-action; so the decision maker faces a familiar trade-off between exploration and exploitation. However, in many settings, observing contextual information is costly so the problem is not only what actions to take/ which decisions to make but also what sources of contextual information to observe. Decision-making in the face of costly observations can be formalized as an ordinary MAB problem by combining the choice of the the context to observe and the decision-action to be taken as a single *meta-action* and folding the costs of observations into rewards. However, that approach is so inefficient as to be impractical for any realistic problem.

This dual learning problem can be formalized as a 2-stage Markov Decision Process (MDP) in which decision maker decides to observe subset of the sources and then takes a decision-action based on the observed contextual information from these sources. We use the ideas of (Ortner & Auer, 2007) to design our algorithms. We provide simpler algorithms and regret bounds that have better scaling with the problem parameters by using the specific structure in this dual-learning problem. These algorithms simultaneously select sources of contextual information and actions based on obtained context to maximize the reward minus cost of obtaining these sources. Our first algorithm is called the Optimistic Policy Selection (OPS) and is designed for scenarios in which no initial contextual information is provided to decision-maker. Our second algorithm, Contextual Optimistic Policy Selection(COPS), is identical to OPS except that decision-maker has initial information that can be used to guide observation-actions.

To illustrate the ubiquitous scenarios where our proposed algorithm are useful, let us consider the problem of treating a sequence of patients. Some information about the patients arrives (essentially) free : age, weight, height, family history etc. However, most of the important information can be only obtained by conducting diagnostic tests. In

Proceedings of the 32^{nd} International Conference on Machine Learning, Lille, France, 2015. JMLR: W&CP volume 37. Copyright 2015 by the author(s).

principle, many diagnostic tests can be applied to a given patient before a treatment decision is made. However, each of these tests impose costs : time, money, inconvenience, discomfort and risk. Hence, an important aspect of the decision making is *which* tests to conduct to inform *what* actions to take. In this setting, our algorithms can be viewed as clinical decision support systems. While in the above medical example, the context represented information that is revealed by taking tests, the considered formalism and the algorithms which we develop can be applied in numerous other settings. For instance, in many Smart Cities applications, the contexts can be audio-visual information collected from various locations in the city, and/or features extracted by processing this information. The cost in the first case can be the communication cost or delay associated with transmitting the information to the decision maker and in the second case can be the processing power associated with processing this information. In summary, the contexts can be raw sources of information and/or features/knowledge extracted from this information and the costs can be related to the acquisition, transmission, processing of this information.

2. Related Work

The work most related to ours is (Szepesvari & Zolghadr, 2013) where the learner tries to learn a single best prediction function when there is a cost associated with observing the features/ sources of contextual information in an adversarial setting. Our work differs from this work because we consider a stochastic setting, only bandit feedback from the selected actions and a context-dependent oracle.

Several other works consider various features of our methods: 1) discovering the relevant information, 2) costly information acquisition, and 3) best decision-action to take given available information. For example, in (Cesa Bianchi et al., 2011; Hazan & Koren, 2012), the decision-maker's goal is to adaptively choose which features of the next training example to train a linear regression model having restricted to access to only a subset of the features extracted from the set of training examples. (Tekin & Van Der Schaar, 2014) learns what is relevant information to consider for each decision. However, the complete information is revealed to the decision maker unlike our setting. Within the field of active sensing, (Yu et al., 2009; Greiner et al., 2002) discover relevant information and make a classic prediction in a batch setting rather than in an online setting where information acquisition is costly and decision-actions are taken.

In (Golovin & Krause, 2010), the decision-maker is constrained to select at most m of the actions to maximize the information gain. This work is extended to the online setting in (Gabillon et al., 2013) where the sources/actions are assumed to independent. Our work differs from this work in various dimensions: in our setting sources of the contextual information can be correlated; there is no feedback about information gain of source is received or known; and the contextual information from the various sources is collected at once. Other related work includes (Rostamizadeh et al., 2011; Dekel et al., 2010) where the features are missing/ corrupted. For example, (Dekel et al., 2010) proposes techniques that are robust to classification time noise. In our setting, the forecaster has the option to purchase/obtain these missing/corrupted features by paying a cost.

3. Problem Formulation

In the considered contextual MAB setting, we assume that the contextual information is taken from a finite set of sources $\mathcal{D} = \{1, 2, \ldots, D\}$. The contextual information yielded by source i at time t is denoted by $x_{i,t} \in \mathcal{X}_i \subseteq \mathbb{R}$. Let $\mathcal{X} \subseteq \mathbb{R}^D$ be the context space. In practice, obtaining the contextual information comes at a cost. The cost of obtaining the context associated with the informational source i is assumed to be $c_i \geq 0$. We also assume that if a decision maker does not pay the cost associated with observing the information of a certain source, then that context is not available to the decision maker and it can not used when taking actions. After obtaining the contextual information, the decision maker chooses a *decision-action* a from a finite set $\mathcal{A} = \{1, 2, \ldots, A\}$ and obtains a reward $r_t(a)$. At each time$t$, the following sequence of events is taking place:

i The environment draws the contextual information $\boldsymbol{x}_t = (x_{1,t}, x_{2,t}, \ldots, x_{D,t})$ and reward vector $\boldsymbol{r}_t = (r_t(1), r_t(2), \ldots, r_t(A))$ is drawn from a joint distribution F, i.e., $(\boldsymbol{x}_t, \boldsymbol{r}_t) \sim F$ with $r_t(a) \in [0, 1]$ for all a. Neither the contextual information $\boldsymbol{x}_t$ nor the reward vector $\boldsymbol{r}_t$ is revealed to the decision-maker. This is a fundamental difference between our formulation and the standard contextual multi-armed bandits in which the entire context vector is revealed to the decision-maker at each time t.

ii The decision-maker decides to observe $\mathcal{O}_t \subseteq \mathcal{P}_{\leq m}(\mathcal{D})$, where $\mathcal{P}_{\leq m}(\mathcal{D})$ is the power set of source set $\mathcal{D}$ with cardinality at most m. The decision-maker observes the contextual information $x_{i,t}$ if the cost c_i associated with source i is paid, i.e., $i \in \mathcal{O}_t$.

iii The decision-maker takes a decision-action $a_t \in \mathcal{A}$ and observes a random reward of $r_t(a_t)$.

We map the contextual information from obtained sources to a state vector $\boldsymbol{s}_t = (s_{1,t}, s_{2,t}, \ldots, s_{D,t}) \in \mathcal{S}$ as

$$s_{i,t} = \begin{cases} \phi_i(x_{i,t}) & \text{if } i \in \mathcal{O}_t \\ \varnothing & \text{if } i \notin \mathcal{O}_t, \end{cases}$$

where $\varnothing$ denotes the unobserved contexts and $\phi_i : \mathcal{X}_i \to \mathcal{S}_i$ is *state mapping function* of source i with $|\mathcal{S}_i| = S_i < \infty$. Let $|\mathcal{S}| = S$. We assume that the decision-maker knows the state mapping $(\phi_i(\cdot))_{i\in\mathcal{D}}$ and thus can compute the state after taking an observation-action $\mathcal{O}$. If $\mathcal{X}_i$ is finite, then ϕ_i can be an *identity mapping*, whereas if $\mathcal{X}_i$ is infinite, then ϕ_i can be a partitioning over $\mathcal{X}_i$ based on similarity information that is known a priori. This is a natural assumption– that is, in many practical applications the results of a measurement are discretized into a finite set.

Let $\mathcal{S}(\mathcal{O})$ denote set of states that are *consistent* with observation-action $\mathcal{O}$ and $\mathrm{dom}(\boldsymbol{s})$ denote the observed contextual information in state $\boldsymbol{s}$. Formally,

$$\mathcal{S}(\mathcal{O}) = \{\boldsymbol{s} \in \mathcal{S} : s_i = \varnothing\ \forall i \notin \mathcal{O}, s_i \neq \varnothing\ \forall i \in \mathcal{O}\} \text{ and}$$
$$\boldsymbol{s} \in \mathcal{S}(\mathcal{O}) \iff \mathrm{dom}(\boldsymbol{s}) = \mathcal{O}.$$

We define $\mathcal{X}(\boldsymbol{s})$ to be the set of context information which can be mapped to state $\boldsymbol{s}$ under state mapping function $(\phi_i(\cdot))_{i\in\mathcal{D}}$, i.e.,

$$\mathcal{X}(\boldsymbol{s}) = \{\boldsymbol{x} \in \mathcal{X} : \phi_i(x_i) = s_i\ \forall i \in \mathrm{dom}(\boldsymbol{s})\}.$$

Based on this, we define $\bar{r}(a, \boldsymbol{s})$ to be the expected reward of action a in state $\boldsymbol{s}$ and $p(\boldsymbol{s}|\mathcal{O})$ to be the transition probability to state $\boldsymbol{s} \in \mathcal{S}(\mathcal{O})$ by taking observation-action $\mathcal{O}$. Formally,

$$\bar{r}(a, \boldsymbol{s}) = \mathbb{E}_{(\boldsymbol{x},\boldsymbol{r})\sim F}\left[r(a)|\boldsymbol{x} \in \mathcal{X}(\boldsymbol{s})\right],$$
$$p(\boldsymbol{s}|\mathcal{O}) = \begin{cases} \Pr_{\boldsymbol{x}\sim F_{\mathcal{X}}}\left(\boldsymbol{x} \in \mathcal{X}(\boldsymbol{s})\right) & \text{if } \mathcal{O} = \mathrm{dom}(\boldsymbol{s}) \\ 0 & \text{if } \mathcal{O} \neq \mathrm{dom}(\boldsymbol{s}) \end{cases}$$

where $F_{\mathcal{X}}$ denotes the marginal distribution on the context space $\mathcal{X}$. We denote $\boldsymbol{P} = [p(\boldsymbol{s}|\mathcal{O})]_{\boldsymbol{s}\in\mathcal{S}(\mathcal{O}),\mathcal{O}\in\mathcal{P}_{\leq m}(\mathcal{D})}$ as the state transition matrix and $\boldsymbol{R} = [\bar{r}(a, \boldsymbol{s})]_{a\in\mathcal{A},\boldsymbol{s}\in\mathcal{S}}$ as the expected reward matrix.

Let $\mathcal{H}_t = \{h_\tau\}_{\tau=1}^{t-1}$ be the observation history of the decision-maker up to time t, where $h_\tau = (\mathcal{O}_\tau, \boldsymbol{s}_\tau, a_\tau, r_\tau(a_\tau))$. A *policy* $\pi_t = \{\mathcal{O}_t, \boldsymbol{a}_t\}$ is a mapping from $\mathcal{H}_t$ to a set of sources $\mathcal{O}$ and a decision-action tuple $\boldsymbol{a} = [a(\boldsymbol{s})]_{\boldsymbol{s}\in\mathcal{S}}$, i.e., $\pi_t : \mathcal{H}_t \to \mathcal{P}_{\leq m}(\mathcal{D}) \times \mathcal{S}^{\mathcal{A}}$. Let $\boldsymbol{\pi} = \{\mathcal{O}, \boldsymbol{a}\}$. The space of all deterministic policies is denoted as Π. We define the expected reward attained by following a policy $\boldsymbol{\pi}$ as

$$\rho(\boldsymbol{\pi}, \boldsymbol{P}, \boldsymbol{R}) = \sum_{\boldsymbol{s}\in\mathcal{S}(\mathcal{O})} p(\boldsymbol{s}|\mathcal{O})\bar{r}(a(\boldsymbol{s}), \boldsymbol{s}) - \sum_{i\in\mathcal{O}} c_i.$$

Without loss of generality, we assume that the decision-maker is allowed to select at most m sources of contextual information in each time, where $1 \leq m \leq D$. The expected reward of an *oracle* that follows the best policy (in terms of the expected reward) can be defined as

$$\begin{aligned} \rho^*_m = \underset{\boldsymbol{\pi}\in\boldsymbol{\Pi}}{\text{maximize}} \quad & \rho(\boldsymbol{\pi}, \boldsymbol{P}, \boldsymbol{R}) \\ \textit{subject to} \quad & \mathcal{O} \in \mathcal{P}_{\leq m}(\mathcal{D}). \end{aligned} \tag{1}$$

We introduce the following notation to define the oracle's policy. Let $F(\boldsymbol{s})$ denote the expected reward of the best action at state $\boldsymbol{s}$, and $V(\mathcal{O})$ denote the value of observation-action $\mathcal{O}$. Formally,

$$F(\boldsymbol{s}) = \max_{a\in\mathcal{A}} \bar{r}(a, \boldsymbol{s}),$$
$$V(\mathcal{O}) = \sum_{\boldsymbol{s}\in\mathcal{S}(\mathcal{O})} p(\boldsymbol{s}|\mathcal{O})F(\boldsymbol{s}) - \sum_{i\in\mathcal{O}} c_i.$$

The oracle policy $\pi^* = \{O^*_m, \boldsymbol{a}^*\}$ is defined as

$$\begin{aligned} a^*(s) &= \arg\max_{a\in\mathcal{A}} \bar{r}(a, \boldsymbol{s}), \\ \mathcal{O}^*_m &= \arg\max_{\mathcal{O}\in\mathcal{P}_{\leq m}(\mathcal{D})} V(\mathcal{O}). \end{aligned} \tag{2}$$

The T-time regret of the policy $\boldsymbol{\pi} = [\pi_t = (\mathcal{O}_t, \boldsymbol{a}_t)]_{t=1}^T$ compared to the oracle's policy is given by

$$\mathrm{Reg}^{\boldsymbol{\pi}}_T \triangleq T\rho^*_m - \sum_{t=1}^{T} \rho(\pi_t, \boldsymbol{P}, \boldsymbol{R}). \tag{3}$$

In the rest of this paper, we will focus on computing the optimal policy which minimizes the regret, i.e.,

$$\begin{aligned} \underset{\boldsymbol{\pi}}{\text{minimize}} \quad & \mathrm{Reg}^{\boldsymbol{\pi}}_T \\ \textit{subject to} \quad & \mathcal{O}_t \in \mathcal{P}_{\leq m}(\mathcal{D}),\ \forall t = 1, 2, \ldots, T. \end{aligned}$$

One solution is to define a set of *meta-actions* that comprises all combinations of observation-action and decision-action selections for each deterministic policy, and then apply a standard MAB algorithm (such as the UCB algorithm (Auer et al., 2002)) while considering these meta-actions to be the action space. While this algorithm is straightforward to implement, it scales with the number of deterministic policies. It can easily be seen that $|\Pi| = |\mathcal{P}_{\leq m}(\mathcal{D})|A^S$. This makes the algorithm computationally infeasible even for a small numbers of actions and states, e.g. $A = 10$ and $S = 10$. This scaling behavior follows from the fact that the algorithm does not account for policies that share the same action, in the sense that selecting an action yields information for many policies. Stemming from this observation and using the theory of undiscounted MDPs (Ortner & Auer, 2007), we construct more efficient algorithms in the next section.

4. Optimistic Policy Selection(OPS) Algorithm

Since the decision-maker does not know the true expected rewards and state transition probabilities, we develop the OPS algorithm to solve the policy search problem in an optimistic way. This can be performed using the following

optimistic policy search:

$$\underset{\boldsymbol{\pi}=\{\mathcal{O},[a(\boldsymbol{s})]_{\boldsymbol{s}\in\mathcal{S}}\}}{\text{maximize}} \sum_{\boldsymbol{s}\in\mathcal{S}(\mathcal{O})} \tilde{p}(\boldsymbol{s}|\mathcal{O})\tilde{r}(a(\boldsymbol{s}),\boldsymbol{s}) - \sum_{i\in\mathcal{O}} c_i$$

subject to

$$\Pr\left(\tilde{r}(a(\boldsymbol{s}),\boldsymbol{s}) \geq \bar{r}(a(\boldsymbol{s}),\boldsymbol{s})\right) \geq 1-\delta_y \;\forall \boldsymbol{s}\in\mathcal{S}(\mathcal{O})$$
$$\Pr\left(|\tilde{p}(\boldsymbol{s}|\mathcal{O}) - p(\boldsymbol{s}|\mathcal{O})| \leq \epsilon_d\right) \geq 1-\delta_d \quad \forall \boldsymbol{s}\in\mathcal{S}(\mathcal{O})$$
$$\sum_{\boldsymbol{s}\in\mathcal{S}(\mathcal{O})} \tilde{p}(\boldsymbol{s}|\mathcal{O}) = 1, \quad \forall \mathcal{O}\in\mathcal{P}_{\leq m}(\mathcal{D}), \tag{4}$$

for some $\epsilon_y, \delta_d, \delta_y > 0$. Eq.(4) is known as *optimism in the face of uncertainty* (Auer et al., 2002; Ortner & Auer, 2007). We reformulate the optimization problem (4) using the sample mean estimates and upper confidence bounds. The decision-maker keeps track of the estimates of the mean rewards and the state transition probabilities. At each time t, we define

$$\mathcal{E}_t(\mathcal{O},\boldsymbol{s}) \triangleq \{0<\tau<t : \mathcal{O}_\tau = \mathcal{O}, \boldsymbol{s}_\tau = \boldsymbol{s}\},$$
$$\mathcal{E}_t(\mathcal{O}) \triangleq \{0<\tau<t : \mathcal{O}_\tau = \mathcal{O}\},$$
$$\mathcal{E}_t(a,\boldsymbol{s}) \triangleq \{0<\tau<t : a_\tau = a, \boldsymbol{s}_\tau = \boldsymbol{s}\}.$$

Let $N_t(\mathcal{O},\boldsymbol{s}) = |\mathcal{E}_t(\mathcal{O},\boldsymbol{s})|$, $N_t(\mathcal{O}) = |\mathcal{E}_t(\mathcal{O})|$, $N_t(a,\boldsymbol{s}) = |\mathcal{E}_t(a,\boldsymbol{s})|$. Let $N_t(\boldsymbol{s}) = \sum_{a\in\mathcal{A}} N_t(a,s)$. Furthermore, we can express the estimated mean rewards and state transition probabilities as follows:

$$\hat{r}_t(a,\boldsymbol{s}) = \frac{1}{N_t(a,\boldsymbol{s})} \sum_{\tau\in\mathcal{E}_t(a,\boldsymbol{s})} r_\tau(a),$$
$$\hat{p}_t(\boldsymbol{s}|\mathcal{O}) = \frac{N_t(\mathcal{O},\boldsymbol{s})}{N_t(\mathcal{O})},$$

provided that $N_t(a,\boldsymbol{s}) > 0$ and $N_t(\mathcal{O}) > 0$. Since these estimates can deviate from their true mean values, we need to add appropriate confidence intervals when searching for the optimal policy.

The OPS operates in rounds $k = 1, 2, \ldots$. In the beginning of each round k, the OPS solves the following optimization problem to find the policy to follow in that round:

$$\underset{\boldsymbol{\pi}=\{\mathcal{O},[a(\boldsymbol{s})]_{\boldsymbol{s}\in\mathcal{S}}\}}{\text{maximize}} \sum_{\boldsymbol{s}\in\mathcal{S}(\mathcal{O})} \tilde{p}(\boldsymbol{s}|\mathcal{O})\tilde{r}(a(\boldsymbol{s}),\boldsymbol{s}) - \sum_{i\in\mathcal{O}} c_i$$

subject to

$$\tilde{r}(a,\boldsymbol{s}) \leq \hat{r}_{t_k}(a,\boldsymbol{s}) + \text{conf}_{1,t_k}(a,\boldsymbol{s}), \quad \forall (a,\boldsymbol{s}) \in (\mathcal{A},\mathcal{S})$$
$$|\tilde{p}(\boldsymbol{s}|\mathcal{O}) - \hat{p}_{t_k}(\boldsymbol{s}|\mathcal{O})| \leq \text{conf}_{2,t_k}(\mathcal{O}), \quad \boldsymbol{s}\in\mathcal{S}(\mathcal{O})$$
$$\sum_{\boldsymbol{s}\in\mathcal{S}(\mathcal{O})} \tilde{p}(\boldsymbol{s}|O) = 1, \quad \mathcal{O}\in\mathcal{P}_{\leq m}(\mathcal{D}), \tag{5}$$

where $\text{conf}_{1,t}(a,\boldsymbol{s})$ and $\text{conf}_{2,t}(\mathcal{O})$ are the confidence bounds on the estimators. The solution to the optimization problem (5) is higher than the reward of oracle solution ρ^*_m with high probability, (see Sec. 8 Lemma 2). Let

Algorithm 1 Optimistic Policy Selection(OPS)

Input: $(\phi_i(\cdot))_{i\in\mathcal{D}}, \alpha, m$
Initialize: $\text{conf}_{1,t}(a,\boldsymbol{s}) = \min\left(1, \sqrt{\frac{\log(2t^\alpha SA)}{2N_t(a,\boldsymbol{s})}}\right)$
Initialize: $\text{conf}_{2,t}(\mathcal{O}) = \min\left(1, \sqrt{\frac{\log(4t^\alpha|\mathcal{S}(\mathcal{O})|)}{2N_t(\mathcal{O})}}\right)$
for rounds $k = 1, 2, \ldots$ **do**
  $t_k = t$
  Solve for $\hat{a}_{t_k}(\boldsymbol{s})$ for all $\boldsymbol{s}\in S$.
  Solve LP given in (5) for all $\mathcal{O}\in\mathcal{P}_{\leq m}(\mathcal{D})$.
  $\hat{\mathcal{O}}_{t_k} = \arg\max_{\mathcal{O}\in\mathcal{P}_{\leq m}(\mathcal{D})} \hat{V}_t(\mathcal{O})$.
  repeat
    Take observation-action $\hat{\mathcal{O}}_{t_k}$
    $s_{i,t} = \phi_i(x_{i,t})$ for all $i\in\hat{\mathcal{O}}_{t_k}$.
    $s_{i,t} = \varnothing$ for all $i\notin\hat{\mathcal{O}}_{t_k}$.
    Take action $\hat{a}_{t_k}(\boldsymbol{s}_t)$.
    Observe reward $r_t(\hat{a}_{t_k}(\boldsymbol{s}_t))$.
  until $\forall \boldsymbol{s}\in\mathcal{S}$ and $a\in\mathcal{A}: N_t(\boldsymbol{s}) < 2N_{t_k}(\boldsymbol{s})$
  $t = t+1$
end for

$\boldsymbol{\pi}_{t_k} = \{\hat{\mathcal{O}}_{t_k}, [\hat{a}_{t_k}(\boldsymbol{s})]_{\boldsymbol{s}\in\mathcal{S}}\}$ be the maximizer of the optimization problem in (5). OPS follows the policy $\boldsymbol{\pi}_{t_k}$ until the visits to one of the states is doubled. In general the optimization problem (5) is nonlinear and nonconvex. As we show (5) can be reduced to a set of Linear Programming (LP) problems which have polynomial time complexity and can be solved efficiently (Boyd & Vandenberghe, 2004). Before presenting the Proposition 1 to reduce (5) into a set of LPs, we introduce some notation. Let

$$\hat{f}_t(a,\boldsymbol{s}) = \hat{r}_t(a,\boldsymbol{s}) + \text{conf}_{1,t}(a,\boldsymbol{s}),$$
$$\hat{a}_t(\boldsymbol{s}) = \arg\max_{a\in\mathcal{A}} \hat{f}_t(a,\boldsymbol{s}),$$
$$\hat{F}_t(\boldsymbol{s}) = \hat{f}_t(\hat{a}_t(\boldsymbol{s}),\boldsymbol{s}). \tag{6}$$

For all $\mathcal{O}\in\mathcal{P}_{\leq m}(\mathcal{D})$ and t_k, we define the optimistic value function of observation-action $\mathcal{O}$ (denoted as $\hat{V}_{t_k}(\mathcal{O})$) as the solution to following LP:

$$\hat{V}_{t_k}(\mathcal{O}) = \underset{[\tilde{p}(\boldsymbol{s}|\mathcal{O})]_{\boldsymbol{s}\in\mathcal{S}(\mathcal{O})}}{\text{maximize}} \sum_{\boldsymbol{s}\in\mathcal{S}(\mathcal{O})} \tilde{p}(\boldsymbol{s}|\mathcal{O})\hat{F}_t(\boldsymbol{s}) - \sum_{i\in\mathcal{O}} c_i$$

subject to

$$|\tilde{p}(\boldsymbol{s}|\mathcal{O}) - \hat{p}_{t_k}(\boldsymbol{s}|\mathcal{O})| \leq \text{conf}_{2,t_k}(\mathcal{O}) \;\forall \boldsymbol{s}\in\mathcal{S}(\mathcal{O})$$
$$\sum_{\boldsymbol{s}\in\mathcal{S}(\mathcal{O})} \tilde{p}(\boldsymbol{s}|\mathcal{O}) = 1. \tag{7}$$

With this notation, we can now present Proposition 1, which shows that the problem in (5) can be converted into a set of LPs.

Proposition 1 *At any time t, the optimization in (5) can be solved using:*

1. $\hat{a}_{t_k}(\boldsymbol{s}) = \arg\max_{a\in\mathcal{A}} \hat{f}_{t_k}(a,\boldsymbol{s})$.

2. $\hat{\mathcal{O}}_{t_k} = \arg\max_{\mathcal{O}\in\mathcal{P}_{\leq m}(\mathcal{D})} \hat{V}_{t_k}(\mathcal{O})$.

Proposition 1 shows that the non-convex optimization problem given in (5) is reduced to a set of LPs for each observation-action $\mathcal{O}$. Based on Proposition 1, in each time t of round k, the OPS observes the sources of contextual information from the set $\hat{\mathcal{O}}_{t_k}$ and takes the decision-action $\hat{a}_{t_k}(\boldsymbol{s}_t)$ in state $\boldsymbol{s}_t$. The pseudo-code of the OPS is given in Algorithm 1.

4.1. Regret Upper Bounds for OPS

Let $\Delta(\mathcal{O})$ be the minimum suboptimality gap among the policies that select observation-action $\mathcal{O}$. Formally, for suboptimal observation-action $\mathcal{O} \in \mathcal{P}_{\leq m}(\mathcal{D}) \setminus \mathcal{O}^*_m$,

$$\Delta(\mathcal{O}) \triangleq V(\mathcal{O}^*_m) - V(\mathcal{O}) \tag{8}$$

where $\mathcal{O}^*_m$ is defined in (2). Recall that the best policy optimization is given by (1), from which it follows that $V(\mathcal{O}^*_m) = \rho^*_m$. With minimal notational abuse, we denote $\Delta(\mathcal{O}^*_m)$ to be the minimum suboptimal gap among the policies that selects optimal observation-action $\mathcal{O}^*_m$. Formally,

$$\begin{aligned}\Delta_{\min}(\boldsymbol{s}) &= \bar{r}(a^*(\boldsymbol{s}), \boldsymbol{s}) - \max_{a\in\mathcal{A}\setminus a^*(\boldsymbol{s})} \bar{r}(a, \boldsymbol{s}),\\ \Delta(\mathcal{O}^*_m) &= \min_{\boldsymbol{s}\in\mathcal{S}(\mathcal{O}^*_m)} p(\boldsymbol{s}|\mathcal{O}^*_m)\Delta_{\min}(\boldsymbol{s}),\end{aligned} \tag{9}$$

where $\Delta_{\min}(\boldsymbol{s})$ is the minimal suboptimality gap at state $\boldsymbol{s}$. Let $\tau(\mathcal{O}) = \max_{\boldsymbol{s}\in\mathcal{S}(\mathcal{O})} 1/p(\boldsymbol{s}|\mathcal{O})$, where we assume $p(\boldsymbol{s}|\mathcal{O}) > 0$ for all $\boldsymbol{s} \in \mathcal{S}(\mathcal{O})$. Note that $p(\boldsymbol{s}|\mathcal{O}) = 0$ means that state $\boldsymbol{s}$ is not observed under observation-action $\mathcal{O}$. Therefore, we can exclude such states from state space $\mathcal{S}(\mathcal{O})$ and, we can assume $\tau(\mathcal{O}) < \infty$ without loss of generality. Theorem 1 bounds the regret of OPS algorithm.

Theorem 1 *If the OPS runs with $\alpha > 2$ until time T, then the regret of the OPS algorithm (1) satisfies:*

$$Reg^{OPS}_T \leq \sum_{\mathcal{O}\in\mathcal{P}_{\leq m}(\mathcal{D})} \left[\frac{C_1 A|\mathcal{S}(\mathcal{O})|^2 \log\left(4SAT^{\gamma}\right)}{\Delta(\mathcal{O})} + 2AS^2\tau(\mathcal{O})\log_2(S^{-1}T)\right] + C_2$$

where $C_1, C_2 > 0$ are constants that are independent of the time horizon T.

The proof of this theorem can be found in Sec. 8. As proven in Theorem 1, the regret of the OPS algorithm scales with $O(|\mathcal{P}_{\leq m}(\mathcal{D})|S^2A)$. Recall from Sec. 3 that the UCB-based algorithm scales with $O(|\mathcal{P}_{\leq m}(\mathcal{D})|A^S)$. Therefore the OPS algorithm provides a significant regret minimization in terms of its dependence on the number of actions and states compared with the UCB-based algorithm. To illustrate the bound, consider a d dimensional context space $\mathcal{X}$ where each context space $\mathcal{X}_i$ is finite with $|\mathcal{X}_i| = M_i$. Let $\phi_i(\cdot)$ be an identity operator and $m = 1$. In this case, $S = 1 + \sum_{i=1}^{D} M_i$ and $|\mathcal{P}_{\leq 1}(\mathcal{D})| = D + 1$. Therefore, the regret of the OPS algorithm scales polynomially in the number of contexts, actions, number of sources as $\mathcal{O}(D^3 M^2_{\max} A)$ where $M_{\max} = \max_i M_i$. In general, the OPS algorithm scales polynomially as $O((D^3 M^2_{\max})^m A)$ when $m << D$. In many real-world scenarios (e.g. transportation, automation, security, medical treatment protocols), the decisions are informed by only a few number of relevant sources.

5. Contextual Optimistic Policy Selection (COPS) Algorithm

When formulating the OPS Algorithm, we assume that no initial sources of contextual information are available to the decision-maker. However, in some applications it may be possible for the decision-maker to obtain contextual information for free. This information can then be exploited to guide the decision-makers selection of additional sources.

In this section we extend the OPS Algorithm to account for this freely available contextual information. Let $\mathcal{X} = \mathcal{X}_0 \times \mathcal{X}_1$ denote a $(d+r)$-dimensional context space where $\mathcal{X}_0$ denotes a r-dimensional free context space and $\mathcal{X}_1$ denotes a d-dimensional context space which is costly to observe. In each time t, the decision-maker observes an initial contextual information $x^0_t \in \mathcal{X}_0$, decides to observe the sources $\mathcal{O}_t$ and the contextual information from these sources is revealed to the decision-maker. Based on all observed contextual information, the decision-maker takes a decision-action a_t and observes a random reward $r_t(a_t)$.

Let $\boldsymbol{s}_{0,t}$ denote the initial state vector before having taken observation-action $\mathcal{O}_t$ and $\boldsymbol{s}_{1,t}$ denote the final state vector after taking observation-action $\mathcal{O}_t$. Formally, $\boldsymbol{s}_{0,t} = \left[\phi_1(x^0_{1,t}), \ldots, \phi_r(x^0_{r,t})\right]$ and $\boldsymbol{s}_{1,t} = [\boldsymbol{s}_{0,t}, \tilde{\boldsymbol{s}}_t]$

$$\tilde{s}_{i,t} = \begin{cases}\phi_{i+r}(x^1_{i,t}) & \text{if } i \in \mathcal{O}_t\\ \varnothing & \text{if } i \notin \mathcal{O}_t\end{cases}$$

To aid in the construction of the COPS algorithm, we introduce the following notation. Let $\mathcal{S}_0$ denote the set of all initial states, and $\mathcal{S}(O, \boldsymbol{s})$ denote the set of all final states with initial state $\boldsymbol{s}$ and that are consistent with observation-action $\mathcal{O}$. Let $\mathcal{S} = \cup_{\mathcal{O}\in\mathcal{P}_{\leq m}(\mathcal{D})} \cup_{\boldsymbol{s}\in\mathcal{S}_0} \mathcal{S}(O, \boldsymbol{s})$ denote the set of all states, where

$$\begin{aligned}\mathcal{S}(\mathcal{O}, \boldsymbol{s}) = \{\tilde{\boldsymbol{s}} \in \mathcal{S} : \tilde{s}_i = s_i \ \forall i \in \mathrm{dom}(\boldsymbol{s}),\\ \tilde{s}_i = \varnothing \ \forall i \notin \mathrm{dom}(\boldsymbol{s}) \cup \mathcal{O}, \tilde{s}_i \neq \varnothing \ \forall i \in \mathcal{O}\}.\end{aligned}$$

Additionally, define $p(\boldsymbol{s}_0)$ as the probability of having an initial state of $\boldsymbol{s}_0$, and $p(\boldsymbol{s}_1|\boldsymbol{s}_0, \mathcal{O})$ as the transition probability from initial state $\boldsymbol{s}_0$ to the final state $\boldsymbol{s}_1$ by taking

observation-action $\mathcal{O}$. The value of the observation-action $\mathcal{O}$ in state $\boldsymbol{s}_0$ is given by $V(\mathcal{O}, \boldsymbol{s}_0)$ with $O^*_m(\boldsymbol{s}_0)$ denoting the optimal observation-action in state $\boldsymbol{s}$, i.e.,

$$V(\mathcal{O}, \boldsymbol{s}_0) = \sum_{\boldsymbol{s}_1 \in S_{\mathcal{O},\boldsymbol{s}_0}} p(\boldsymbol{s}_1|\boldsymbol{s}_0, \mathcal{O}) F(\boldsymbol{s}_1) - \sum_{i \in \mathcal{O}} c_i,$$
$$\mathcal{O}^*_m(\boldsymbol{s}_0) = \arg\max_{\mathcal{O} \in \mathcal{P}_{\leq m}(\mathcal{D})} V(\mathcal{O}, \boldsymbol{s}_0).$$

With these parameters defined, a policy $\boldsymbol{\pi}_t$ is a mapping from the history $\mathcal{H}_t$ to a tuple of observation-action and decision-action. Formally, $\boldsymbol{\pi} = \{\boldsymbol{O}, \boldsymbol{a}\}$ where $\boldsymbol{O} = [\mathcal{O}(\boldsymbol{s}_0)]_{\boldsymbol{s}_0 \in \mathcal{S}_0}$. The expected reward of the oracle policy $\boldsymbol{\pi}^* = \left\{ [O^*_m(\boldsymbol{s}_0)]_{\boldsymbol{s}_0 \in \mathcal{S}_0}, [a^*(\boldsymbol{s})]_{\boldsymbol{s} \in \mathcal{S}} \right\}$ is given by:

$$\rho^*_m = \sum_{\boldsymbol{s}_0 \in S_0} p(\boldsymbol{s}_0) \left[V(\mathcal{O}^*_m(\boldsymbol{s}_0), \boldsymbol{s}_0) - \sum_{i \in \mathcal{O}^*_m(\boldsymbol{s}_0)} c_i \right].$$

The modified OPS algorithm is denoted as the Contextual Optimistic Policy Selection (COPS) algorithm as it accounts for the initial contextual information when taking observation-actions. Let $\hat{p}_t(\boldsymbol{s}_1|\boldsymbol{s}_0, \mathcal{O})$ be the state transition estimate from initial state $\boldsymbol{s}$ to the final state $\boldsymbol{s}_1 \in \mathcal{S}(\mathcal{O}, \boldsymbol{s}_0)$ when observation-action $\mathcal{O}$ is taken. In the beginning of round k, the COPS first computes the optimistic rewards of state-action pairs and $\hat{F}_{t_k}(\boldsymbol{s})$ defined in (6). Based on the evaluated results, the COPS algorithm then computes $\hat{V}_t(\mathcal{O}, \boldsymbol{s}_0)$ for all $\boldsymbol{s}_0 \in \mathcal{S}_0$ by solving the following LP problem:

$$\hat{V}_t(\mathcal{O}, \boldsymbol{s}_0)$$
$$= \underset{[\tilde{p}(\boldsymbol{s}_1|\boldsymbol{s}_0,\mathcal{O})]_{\boldsymbol{s}_1 \in \mathcal{S}(\mathcal{O},\boldsymbol{s}_0)}}{\text{maximize}} \sum_{\boldsymbol{s}_1 \in \mathcal{S}(\mathcal{O},\boldsymbol{s}_0)} \tilde{p}(\boldsymbol{s}_1|\boldsymbol{s}_0, \mathcal{O}) \hat{F}_t(\boldsymbol{s}_0) - \sum_{i \in O} c_i$$

subject to

$$|\tilde{p}(\boldsymbol{s}_1|\boldsymbol{s}_0, \mathcal{O}) - \hat{p}_t(\boldsymbol{s}_1|\boldsymbol{s}_0, \mathcal{O})| \leq \text{conf}_{2,t}(\mathcal{O}, \boldsymbol{s}_0),$$
$$\sum_{\boldsymbol{s}_1 \in \mathcal{S}(\mathcal{O},\boldsymbol{s}_0)} \tilde{p}(\boldsymbol{s}_1|\boldsymbol{s}_0, \mathcal{O}) = 1 \; \forall \boldsymbol{s}_1 \in \mathcal{S}(\mathcal{O}, \boldsymbol{s}_0), \tag{10}$$

where maximization is over parameters $[\tilde{p}(\boldsymbol{s}_1|\boldsymbol{s}_0, \mathcal{O})]_{\boldsymbol{s}_1 \in \mathcal{S}(\mathcal{O},\boldsymbol{s}_0)}$. In time t of round k, the COPS algorithm takes an observation-action with the highest $\hat{V}_{t_k}(\mathcal{O}, \boldsymbol{s}_{0,t})$ based on the initial state $\boldsymbol{s}_{0,t}$. Then, the algorithm observes the obtained contextual information from these sources and takes a decision-action a_t with the highest $\hat{f}_{t_k}(a, \boldsymbol{s}_{1,t})$ based on the final state $\boldsymbol{s}_{1,t}$.

Similar to OPS, we define $\Delta(\mathcal{O}, \boldsymbol{s}_0)$ to be the suboptimality gap of observation-action $\mathcal{O}$. Formally,

$$\Delta(\mathcal{O}, \boldsymbol{s}_0) = V(\mathcal{O}^*_m(\boldsymbol{s}_0), \boldsymbol{s}_0) - V(\mathcal{O}, \boldsymbol{s}_0),$$
$$\Delta(\mathcal{O}^*_m(\boldsymbol{s}_0), \boldsymbol{s}_0) = \min_{\boldsymbol{s}_1 \in S_{\mathcal{O},\boldsymbol{s}_0}} p(\boldsymbol{s}_1|\boldsymbol{s}_0, \mathcal{O}^*_m(s_0)) \Delta_{\min}(\boldsymbol{s}_1),$$

where $\Delta_{\min}(\boldsymbol{s}_1)$ is the minimum suboptimality gap of state $\boldsymbol{s}_1$ defined in Eq.(9). Let $\tau(\mathcal{O}, \boldsymbol{s}_0) = \min_{\boldsymbol{s}_1 \in \mathcal{S}(\mathcal{O},\boldsymbol{s}_0)} 1/p(\boldsymbol{s}_1|\boldsymbol{s}_0, \mathcal{O})$. Using the defined parameters, Theorem 2 provides the regret bound for the COPS algorithm.

Theorem 2 *If the COPS algorithm runs with $\alpha > 2$ until time T, the regret of COPS satisfies*

$$Reg^{C\text{-}OPS}_T$$
$$\leq \sum_{\boldsymbol{s}_0 \in \mathcal{S}_0} \sum_{O \in \mathcal{P}_{\leq m}(\mathcal{D})} p(\boldsymbol{s}_0) \left[\frac{C_3 A |S_{O,\boldsymbol{s}_0}|^2 \log(4SAT^\alpha)}{\Delta(\mathcal{O}, \boldsymbol{s}_0)} + 2AS^2 \tau(O, \boldsymbol{s}_0) \log_2(S^{-1}T) \right] + C_4$$

where $C_3, C_4 > 0$ are independent of time horizon T.

As can be seen from Theorem 2, there is a negligible difference between the OPS and COPS algorithm : they are both having a logarithmic regret. However, the COPS algorithm is expected to achieve higher expected rewards since it exploits some initial contextual information for free.

6. Illustrative Results

In this section we illustrate the performance of the OPS algorithm 1 in a clinical application setting where the goal is to determine which chemotherapy should be given to which patient. The dataset consists of over 300,000 breast cancer patient cases characterized by 4 different contexts (e.g. age, estrogen receptor, tumor state, previous treatments). The data is collected from 80 clinical studies which focus on 2 different chemotherapy treatments (AC and ACT chemotherapies). We define the reward of a treatment (i.e. action) using the probability of successfully curing patients, (i.e. no recurrence of cancer in 5 years).

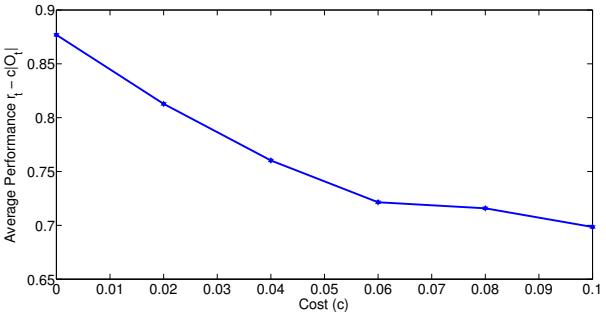

Figure 1. Performance of the OPS with different costs with $m = 4$ at time $T = 300,000$.

Performance of the OPS Algorithm with Different Costs: We consider that for each source $i \in \mathcal{D}$, the cost of observing the associated context $c_i = c$. We illustrate how the results of the OPS Algorithm are dependent on the value of this cost c in Fig. 1. Recall that the clinician

aims to maximize $r_t(a_t) - c|\mathcal{O}_t|$, i.e., the rewards minus cost of observation, where the observation-action is given by $\mathcal{O}_t$ and and the associated treatment is a_t. As Fig. 1 illustrates, the average gain of the OPS algorithm decreases as the costs associated with observing the contextual information increases. However, it should be noted that the OPS algorithm optimally learns the best contextual information to select while simultaneously taking actions irrespective of the costs of observation. Table 1 shows that when the cost is low, the OPS performs better than the UCB by obtaining some of the sources of the contextual information that are well-worth their costs. As the cost increases, the OPS behaves similar to UCB by selecting no sources of contextual information. Therefore, the slope of the performance-cost curve of the OPS illustrated in 1 decreases as the cost of obtaining the sources increases.

-	$c = 0.001$	$c = 0.02$	$c = 0.04$
UCB	79.93%	79.93%	79.93%.
OPS	87.69%	81.28%	76.03%.

Table 1. Performance of the OPS and UCB with different costs with $m = 4$ at time $T = 300,000$.

Performance of the OPS with different m: Consider the situation in which a decision-maker is attempting to decide if observing a single context (i.e. $m = 1$) is necessary. We know from Theorem 1 that the complexity of the OPS algorithm is $O(D^3 M_{\max}^2 k^2)$. However if the decision-maker would like to select from several possible sources $m > 1$, then we have the following results: first, the regret bound increases with m as a result of the possible sources to select; second, the expected reward obtained by the oracle policy increases with m. This results in an increase in the time to learn the optimal sources to select while simultaneously learning to choose the optimal actions. This is a ubiquitous and fundamental problem in any real-world application : as you increase the number of possible observations increase, it becomes more difficult to learn the relationship between the contexts and the rewards. However, we show next that the OPS algorithm can still efficiently account for multiple sources of information in an online framework.

Consider a scenario where patients that have 4 unknown contexts. If $m = 1$ then the OPS algorithm can efficiently learn the best source of contextual information and actions to take to maximize the success of a treatment using a dataset with $T = 300,000$ samples as illustrated in Table 2. However, for the same size T, if we increase the value of the number of possible sources to $m = 2$ to $m = 4$, it is expected the the performance of the OPS algorithm will decrease with increasing m. Remarkably, as Table 2 illustrates, only a very slight decrease in the probability of success for this increase in m. As the results in Table 2 shows, the results of the OPS algorithm naturally illustrate that taking fewer samples may result in a better performance in terms of the rewards.

-	$m = 1$	$m = 2$	$m = 3$	$m = 4$
OPS	91.10%	90.63%	90.23%	89.73%.

Table 2. Performance of the OPS with different m when cost $c_i = c = 0$ at time $T = 300,000$

Performance of OPS versus COPS: We compared the performance of COPS with the setting in which age is observed for free and OPS. In this case, the COPS outperforms the OPS by 1.4% (OPS : 88.6% and COPS : 90%).

7. Conclusion and Future Work

In this paper two reinforcement learning algorithms are presented which simultaneously discover relevant costly sources while learning the best decision-actions to make with the obtained contextual information to maximize the reward in a dynamic setting. The Optimistic Policy Selection (OPS) algorithm is useful when no initial information is available to a decision-maker. If some information is available, then the Contextual Optimistic Policy Selection (COPS) algorithm can utilize it to improve performance. We prove that both algorithms satisfy a regret bound that is logarithmic in time.

Future work of the current algorithms including imposing functional constraints on the expected rewards including bounded variation, Lipschitz smoothness, and linearity. Imposing these structural constraints on the rewards can reduce the complexity of learning the relevant sources.

8. Proofs

In this section, we provide a sketch of the proof of Theorem 1. Theorem 2 is the proof of Theorem 1 repeated for each $\boldsymbol{s}_0 \in \mathcal{S}_0$. The proofs of the lemmas and Proposition 1 are in the supplementary material.

Lemma 1 *Number of rounds at time T is bounded by* $S \log_2 \left(S^{-1} T\right)$.

Let $[\hat{p}_t^*(\boldsymbol{s}|\mathcal{O})]_{\boldsymbol{s}\in\mathcal{S}(\mathcal{O})}$ be the maximizer of (7) for each observation-action $\mathcal{O}$. Let $\hat{\rho}^*(t_k)$ be the solution of (5). Observe that

$$\begin{aligned}\hat{\rho}^*(t_k) &= \arg\max_{\mathcal{O}\in\mathcal{P}_{\leq m}(\mathcal{D})} \hat{V}_{t_k}(\mathcal{O}) \\ &= \arg\max_{\mathcal{O}\in\mathcal{P}_{\leq m}(\mathcal{D})} \sum_{\boldsymbol{s}\in\mathcal{S}(\mathcal{O})} \hat{p}_{t_k}^*(\boldsymbol{s}|\mathcal{O}) \hat{F}_{t_k}(\boldsymbol{s}) - \sum_{i\in\mathcal{O}} c_i\end{aligned}$$

Lemma 2 $\Pr\left(\rho_m^* \geq \hat{\rho}^*(t_k)\right) \leq t_k^{-\alpha}$.

Let $\hat{\boldsymbol{\pi}}_k = \left[\hat{\mathcal{O}}_{t_k}, [\hat{a}_{t_k}(\boldsymbol{s})]_{\boldsymbol{s}\in S}\right]$ be the suboptimal policy computed by the OPS at the beginning of round k. Let Δ_k be the suboptimality gap of policy $\hat{\boldsymbol{\pi}}_k$. Then,

$$\begin{aligned}\Delta_k &= \rho^*_m - \rho(\hat{\boldsymbol{\pi}}_k, \boldsymbol{P}, \boldsymbol{R}) \\ &\overset{\text{w.h.p}}{\leq} \hat{\rho}^*(t_k) - \rho(\hat{\boldsymbol{\pi}}_k, \boldsymbol{P}, \boldsymbol{R})\end{aligned}$$

where inequality follows with probability $1 - t_k^{-\alpha}$ by Lemma 2. If $\hat{\rho}^*(t_k) - \rho(\hat{\boldsymbol{\pi}}_k, \boldsymbol{P}, \boldsymbol{R}) < \Delta_k$, then the policy $\hat{\boldsymbol{\pi}}_k$ can not be computed by the OPS with high probability.

Lemma 3 *Let $\boldsymbol{\pi} = \{\mathcal{O}, \boldsymbol{a}\}$ be an arbitrary policy and $\tilde{\boldsymbol{P}} = [p(\boldsymbol{s}|\mathcal{O})]_{\boldsymbol{s}\in\mathcal{S}(\mathcal{O}),\mathcal{O}\in\mathcal{P}_{\leq m}(\mathcal{D})}$ and $\tilde{\boldsymbol{R}} = [r(a,\boldsymbol{s})]_{a\in A,\boldsymbol{s}\in S}$ be arbitrary transition and expected reward matrix. We have $\rho(\boldsymbol{\pi}, \tilde{\boldsymbol{P}}, \tilde{\boldsymbol{R}}) - \rho(\boldsymbol{\pi}, \boldsymbol{P}, \boldsymbol{R}) < \Delta$ if for all $\boldsymbol{s} \in \mathcal{S}(\mathcal{O})$ it holds that*

$$\begin{aligned}&|\tilde{p}(\boldsymbol{s}|\mathcal{O}) - p(\boldsymbol{s}|\mathcal{O})| < 0.25|\mathcal{S}(\mathcal{O})|^{-1}\Delta \\ &|\tilde{r}(a(\boldsymbol{s}),\boldsymbol{s}) - \bar{r}(a(\boldsymbol{s}),\boldsymbol{s})| < 0.5|\mathcal{S}(\mathcal{O})|^{-1}p^{-1}(\boldsymbol{s}|\mathcal{O})\Delta\end{aligned}$$

Lemma 4
a) If $N_t(a,\boldsymbol{s}) > 8|\mathcal{S}(\mathcal{O})|^2p^2(\boldsymbol{s}|\mathcal{O})\Delta^{-2}\log(2SAt^\alpha)$, then with probability $1 - 0.5|S|^{-1}k^{-1}t^{-\alpha}$

$$|\hat{f}_t(a,\boldsymbol{s}) - \bar{r}(a,\boldsymbol{s})| < 0.5|\mathcal{S}(\mathcal{O})|^{-1}p^{-1}(\boldsymbol{s}|\mathcal{O})\Delta.$$

b) If $N_t(O) > 32|\mathcal{S}(\mathcal{O}|^2\Delta^{-2}\log(4St^\alpha)$, then with probability $1 - 0.5|\mathcal{S}(\mathcal{O})|t^{-\alpha}$

$$|\hat{p}^*_t(\boldsymbol{s}|\mathcal{O}) - p(\boldsymbol{s}|\mathcal{O})| < 0.25|\mathcal{S}(\mathcal{O})|^{-1}\Delta,$$

for all $\boldsymbol{s} \in \mathcal{S}(\mathcal{O})$.

A suboptimal policy $\hat{\boldsymbol{\pi}}_k$ with suboptimality gap Δ_k can not be computed by the OPS with high probability if the following events happen :
a)$N_t(O) > 32|\mathcal{S}(\mathcal{O})|^2\Delta^{-2}\log(4St^\alpha)$
b) There exists state $\boldsymbol{s}$ and decision-action a with $N_t(a,\boldsymbol{s}) > 8|\mathcal{S}(\mathcal{O})|^2p^2(\boldsymbol{s}|\mathcal{O})\Delta^{-2}\log(2SAt^\alpha)$.

Regret due to suboptimal rounds : Suppose that a suboptimal policy $\hat{\boldsymbol{\pi}}_k = (\hat{O}_{t_k}, \hat{\boldsymbol{a}}_{t_k})$ with suboptimality gap Δ_k is computed by the OPS at the beginning of round k.

Consider all suboptimal rounds with $N_t(\mathcal{O}) < 32|\mathcal{S}(\mathcal{O})|^2\Delta_k^{-2}\log(4St^\alpha)$. Let $m(\mathcal{O})$ be the number of such rounds and $\tau_k(\mathcal{O})$ $(k = 1, 2, \ldots, m(\mathcal{O}))$ be their lengths. Then,

$$\sum_{k=1}^{m(\mathcal{O})-1} \tau_k(\mathcal{O}) < 32|\mathcal{S}(\mathcal{O})|^2\Delta_k^{-2}\log(4St^\alpha)$$

Now consider the last round with $\hat{\mathcal{O}}_{t_k} = \mathcal{O}$ which lasts until one of the visits of $(\boldsymbol{s})$ is doubled. Since $\sum_{\boldsymbol{s}\in\mathcal{S}(\mathcal{O})} N_t(\boldsymbol{s}) = N_t(\mathcal{O})$, one of the visits of $(\boldsymbol{s})$ is doubled when $N_t(\mathcal{O})$ is doubled. Therefore,

$$\sum_{k=1}^{m(\mathcal{O})} \tau_k(\mathcal{O}) < 64|\mathcal{S}(\mathcal{O})|^2\Delta_k^{-2}\log(4St^\alpha)$$

Consider all suboptimal rounds with $N_t(a,\boldsymbol{s}) < 8|\mathcal{S}(\mathcal{O})|^2p^2(\boldsymbol{s}|\mathcal{O})\Delta_k^{-2}\log(2SAt^\alpha)$ and $\hat{a}_t(\boldsymbol{s}) = a$. Let $m(a,\boldsymbol{s})$ be the number of these rounds and $\tau_k(a,\boldsymbol{s})$ $(k = 1, \ldots, m(a,s))$ be their lengths. The mean time to observe the state $\boldsymbol{s} \in \mathcal{S}(\mathcal{O})$ under observation-action $\mathcal{O}$ is $1/p(\boldsymbol{s}|\mathcal{O})$. We can divide each round k into $\lfloor p(\boldsymbol{s}|\mathcal{O})\tau_k(a,\boldsymbol{s})/2\rfloor$ steps. By application of Markov inequality, reaching a state $\boldsymbol{s}$ within one step is greater than 0.5. With abuse of notation, let $N_n(a,\boldsymbol{s})$ denote the number of $(\boldsymbol{s},a)$ visits within n steps. Then, by application of Chernoff-Hoeffding inequality,

$$\Pr\left(N_n(a,\boldsymbol{s}) \geq \frac{n}{2} - \sqrt{n\log T}\right) \geq 1 - \frac{1}{T}.$$

Since $N_t(a,\boldsymbol{s}) < 8|\mathcal{S}(\mathcal{O})|^2p^2(\boldsymbol{s}|\mathcal{O})\Delta_k^{-2}\log(2SAt^\alpha)$, with high probability number of such steps satisfy :

$$16|\mathcal{S}(\mathcal{O})|^2p^2(\boldsymbol{s}|\mathcal{O})\Delta_k^{-2}\log(2SAT^\alpha) \geq \frac{n}{2} - \sqrt{n\log T},$$

which implies that number of such steps is less than $c\Delta_k^{-2}|\mathcal{S}(\mathcal{O}|^2p^2(\boldsymbol{s}|\mathcal{O})\log(2T^\alpha SA)$ for some $c < 83$. In other words, with high probability,

$$\begin{aligned}\sum_{k=1}^{m(a,\boldsymbol{s})} &\lfloor p(\boldsymbol{s}|\mathcal{O})\tau_k(a,\boldsymbol{s})/2\rfloor \\ &\leq c\Delta_k^{-2}|\mathcal{S}(\mathcal{O})|^2p^2(\boldsymbol{s}|\mathcal{O})\log(2T^\alpha SA).\end{aligned}$$

Note that $m(\boldsymbol{s},a) \leq |S|A\log_2(S^{-1}A^{-1}T)$ by Lemma 1. Then,

$$\begin{aligned}\mathbb{E}\left(\sum_{k=1}^{m(\boldsymbol{s},a)} \tau_k(\boldsymbol{s},a)\right) &\leq 2p^{-1}(\boldsymbol{s}|\mathcal{O})m(\boldsymbol{s},a) \\ &+ c\Delta_k^{-2}|\mathcal{S}(\mathcal{O})|^2p^2(\boldsymbol{s}|\mathcal{O})\log(2T^\gamma SA) + \frac{1}{T}T.\end{aligned}$$

The regret due to suboptimal policy $\hat{\boldsymbol{\pi}}_k$ is bounded by

$$\begin{aligned}&64|\mathcal{S}(\mathcal{O})|^2\Delta_k\log(4ST^\alpha) \\ &+ c\Delta_k^{-1}|\mathcal{S}(\mathcal{O})|^2p^2(\boldsymbol{s}|\mathcal{O})\log(2T^\gamma SA) \\ &+ 2\Delta_k p^{-1}(\boldsymbol{s}|\mathcal{O})S\log_2(S^{-1}T) + 1\end{aligned}$$

Then, the regret due to any suboptimal observation-action O can be found by summing this up for all state-action pairs and then taking the maximum among the policies with

$\hat{O}_{t_k} = O$. This concludes that regret due to suboptimal observation-action O is bounded by

$$\frac{2c|\mathcal{S}(\mathcal{O})|^2 A \log(4T^\gamma SA)}{\Delta(\mathcal{O})} + 2\tau(\mathcal{O})S^2 \log_2\left(S^{-1}T\right) + SA$$

Regret due to Failure of Confidence Intervals : According to Lemmas and union bound, the probability of failing confidence intervals at round k is bounded by $|\mathcal{P}_{\leq m}(\mathcal{D})|t_k^{-\alpha}$. Let τ_k is the length of the round k, i.e., $\tau_k = t_{k+1} - t_k$. Note that there exists a constant c such that $\tau_k \leq ct_k$. Then,

$$\sum_{k=1}^{\infty} t_k^{-\alpha}\tau_k \leq |\mathcal{P}_{\leq m}(\mathcal{D})| \sum_{k=1}^{\infty} t_k^{-\alpha+1} \leq c^{'}$$

where $c^{'} > 0$ is a constant for $\alpha > 2$.

Summing up the regret due suboptimal policies with observation-action $\mathcal{O}$ and the regret due to failure of confidence intervals concludes the result given in Theorem 1.

References

Auer, Peter, Cesa-Bianchi, Nicolo, and Fischer, Paul. Finite-time analysis of the multi-armed bandit problem. *Machine Learning*, 47:235–256, 2002.

Boyd, Stephen and Vandenberghe, Lieven. *Convex optimization*. Cambridge university press, 2004.

Cesa Bianchi, Nicolo, Shalev Shwartz, Shai, and Shamir, Ohad. Efficient learning with partially observed attributes. *The Journal of Machine Learning Research*, 12:2857–2878, 2011.

Chu, Wei, Li, Lihong, Reyzin, Lev, and Schapire, Robert E. Contextual bandits with linear payoff functions. In *International Conference on Artificial Intelligence and Statistics*, pp. 208–214, 2011.

Dekel, Ofer, Shamir, Ohad, and Xiao, Lin. Learning to classify with missing and corrupted features. *Machine learning*, 81(2):149–178, 2010.

Dudik, Miroslav, Hsu, Daniel, Kale, Satyen, Karampatziakis, Nikos, Langford, John, Reyzin, Lev, and Zhang, Tong. Efficient optimal learning for contextual bandits. *arXiv preprint arXiv:1106.2369*, 2011.

Gabillon, Victor, Kveton, Branislav, Wen, Zheng, Eriksson, Brian, and Muthukrishnan, S. Adaptive submodular maximization in bandit setting. In *Advances in Neural Information Processing Systems*, pp. 2697–2705, 2013.

Golovin, Daniel and Krause, Andreas. Adaptive submodularity: A new approach to active learning and stochastic optimization. In *COLT*, pp. 333–345, 2010.

Greiner, Russell, Grove, Adam J, and Roth, Dan. Learning cost-sensitive active classifiers. *Artificial Intelligence*, 139(2):137–174, 2002.

Hazan, Elad and Koren, Tomer. Linear regression with limited observation. In *Proc. 29th Int. Conf. on Machine Learning*, pp. 807–814, 2012.

Langford, John and Zhang, Tong. The epoch-greedy algorithm for contextual multi-armed bandits. *Advances in Neural Information Processing Systems (NIPS)*, 20: 1096–1103, 2007.

Lu, Tyler, Pál, Dávid, and Pál, Martin. Contextual multi-armed bandits. In *International Conference on Artificial Intelligence and Statistics*, pp. 485–492, 2010.

Ortner, P and Auer, R. Logarithmic online regret bounds for undiscounted reinforcement learning. In *Advances in Neural Information Processing Systems*, 2007.

Rostamizadeh, Afshin, Agarwal, Alekh, and Bartlett, Peter L. Learning with missing features. In *UAI*, pp. 635–642, 2011.

Slivkins, Aleksandrs. Contextual bandits with similarity information. In *24th Annual Conference On Learning Theory*, 2011.

Szepesvari, Csaba and Zolghadr, Navid. Online learning with costly features and labels. In *Advances in Neural Information Processing Systems (NIPS)*, 2013.

Tekin, Cem and Van Der Schaar, Mihaela. Discovering, learning and exploiting relevance. In *Advances in Neural Information Processing Systems*, pp. 1233–1241, 2014.

Yu, Shipeng, Krishnapuram, Balaji, Rosales, Romer, and Rao, R Bharat. Active sensing. In *International Conference on Artificial Intelligence and Statistics*, pp. 639–646, 2009.